

ATTORNEY'S DOCKET NO. IRB-002

PATENTS

UNITED STATES PATENT APPLICATION

OF

ROSS W. CALLON AND THIERRY MAISON

FOR

SYSTEM AND METHOD FOR FACILITATING RECOVERY FROM COMMUNICATION LINK

FAILURES IN A DIGITAL DATA NETWORK

Certificate of Express Mailing

Express Mail Mailing Label No. EI 737 669 115 US

Date of Deposit July 10, 1998

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office To Addressee" Service under 37 CFR 1.10 on the date indicated above and is addressed to the Commissioner of Patents and Trademarks, Washington, D. C. 20231.

By Richard Jordan

Richard A. Jordan

1

FIELD OF THE INVENTION

2 The invention relates generally to the field of digital communications, and more particularly
3 to systems and methods for facilitating recovery from communication link failures in a digital data
4 network.

5

BACKGROUND OF THE INVENTION

6 Digital networks have been developed to facilitate the transfer of information, including data
7 and programs, among digital computer systems and numerous other types of devices. A variety of
8 types of networks have been developed and implemented using diverse information transfer
9 methodologies. In modern networks, information is transferred through a mesh of switching nodes
10 which are interconnected by communication links in a variety of patterns. The mesh interconnection
11 pattern can allow for a number of paths to be available through the network from each computer
12 system or other device which may transmit information as a source device, to another computer
13 system or other device, which is to receive the information as a destination device, so that if
14 congestion develops in particular regions of the network, or if components of the network become
15 inoperative, information can be routed around the congested or inoperative portion of the network.

16 Information transferred from a source device to a destination device is generally transferred
17 in the form of fixed- or variable-length packets, which are received by a switching node over a
18 communication link connected thereto, and transmitted over another communication link to facilitate
19 transfer of the packet to the destination device or another switching node along a path to the
20 destination device. Each packet typically includes address information, including a source address
21 that identifies the particular device that generated the packet, and a destination address that identifies
22 the particular device or devices to receive the packet.

23 Information transferred from a source device to a destination device is generally transferred

1 in the form of fixed- or variable-length packets, which are transferred through the network in
2 accordance with one of several general packet transfer methodologies. In one packet transfer
3 methodology, termed a "connectionless" data transfer methodology, each information packet that
4 is transmitted by a source device to a destination device includes, in addition to the information to
5 be transferred, destination address information which identifies the particular device that is to receive
6 the packet. Each switching node in the network is provided with a database that stores network
7 topology information, which describes the entire topology of the network, and a routing table that
8 provides, among other information, routing information identifying the path to be taken from the
9 switching node to any other switching node in the network. When a switching node receives a
10 packet that is to be transferred to a particular destination device, it (that is, the switching node) will
11 use the routing table to identify a communication link over which the packet is to be transmitted.

12 In another "connection-oriented" packet transfer methodology, packets are transferred
13 through the network's switching nodes over constructs which are generally termed "virtual circuits,"
14 "virtual connections," "switched paths," and the like (generally, "virtual circuits"). When a source
15 device wishes to transfer information to a destination device, initially a preliminary operation will
16 be performed to establish a virtual circuit over communication links defining a path from the source
17 device, through one or more switching nodes to the destination device. In this methodology, each
18 switching node that operates using the virtual circuit methodology includes a virtual circuit table that
19 identifies, for each virtual circuit for which the switching node forms a path, the particular
20 communication link over which the packet is to be transmitted. In addition, each switching node
21 may also include a network topology database that stores network topology information, which may
22 be used for transferring connectionless packets which may be used for a variety of purposes,
23 including, for example, transferring network management information to respective switching nodes.

24 Several problems can arise if a communication link interconnecting two switching nodes fails
25 or otherwise malfunctions. When that occurs, typically some time is required to provide information
26 to all of the switching nodes in the network indicating that the communication link is no longer

1 available. For some time after the failure or malfunction occurs, the network topology information
2 as maintained by the various switching nodes will be inconsistent, in which case their routing
3 determinations, or virtual circuit path establishment decisions will also be inconsistent. In addition,
4 in a connection-oriented network, switching nodes may attempt to establish virtual circuits over the
5 failed communication link. A number of undesirable consequences can arise if switching nodes in
6 a network have inconsistent or incorrect routing information. For example, in a connectionless
7 network, inconsistent or incorrect routing information can result in looping packets among a plurality
8 of switching nodes in the network, loss of packets, as well as routing of packets over paths in the
9 network which are not optimal. In a connection oriented network, inconsistent or incorrect routing
10 information can result in inability to set up new virtual circuits and inability to repair virtual circuits
11 which have failed.

12 SUMMARY OF THE INVENTION

13 The invention provides a new and improved system and method for facilitating recovery from
14 failures of communication links in a digital data network.

15 In brief summary, the invention provides a network comprising a plurality of switching nodes
16 interconnected by communication links for transferring digital packets. At least one switching node
17 in the network pre-establishes a bypass virtual circuit through the network to bypass an element of
18 the network, such as a switching node or a communication link, in the network. The bypass virtual
19 circuit defines a path to another switching node in the network. The first switching node uses the
20 bypass virtual circuit so constructed in forwarding of a packet in the event of a failure or other
21 malfunction of the element if the first switching node would otherwise transfer the packet over that
22 element. By providing and making use of such bypass virtual circuits, the switching nodes which
23 are connected to an element which has failed or otherwise malfunctioned can avoid notifying the
24 other switching nodes, thereby avoiding undesirable consequences that can occur if the routing
25 information maintained by the switching nodes is inconsistent.

BRIEF DESCRIPTION OF THE DRAWINGS

This invention is pointed out with particularity in the appended claims. The above and further advantages of this invention may be better understood by referring to the following description taken in conjunction with the accompanying drawings, in which:

FIG. 1 schematically depicts a computer network including a facility for recovering from communication link failures, constructed in accordance with the invention;

FIG. 2 is a flowchart depicting operations performed by switching nodes in connection with recovering from communication link failures.

DETAILED DESCRIPTION OF AN ILLUSTRATIVE EMBODIMENT

FIG. 1 schematically depicts a computer network 10 including a facility for recovering from communication link failures constructed in accordance with the invention. The computer network 10 includes a plurality of switching nodes 11(1) through 11(N) (generally identified by reference numeral 11(n)) for transferring signals representing data among a number of devices, which in FIG. 1 are represented by packet source/destination devices 12(1) through 12(M) (generally identified by reference numeral 12(m)) in a wide area network ("WAN"). The switching nodes 11(n) used in one embodiment of the invention may be similar to the switching node described in U. S. Patent Application Ser. No. 09/108,771, filed July 2, 1998, in the names of Steven J. Schwartz, et al., entitled "System And Method For Switching Packets In A Network" (Atty. Docket No. TRB-001), assigned to the assignee of the present application and incorporated herein by reference. The packet source/destination devices 12(m), as is conventional, include a particular device, such as a computer system or other device that stores, generates, processes or otherwise uses digital data, a local area network of such devices, or the like (not separately shown) to the wide area network 10. Each packet source/destination device 12(m) is connected over a communication link,

1 generally identified by reference numeral 13(p), to a switching node 11(n) to facilitate transmission
2 of data thereto or the reception of data therefrom. The switching nodes 11(n) are interconnected by
3 communication links, also generally identified by reference numeral 13(p) to facilitate the transfer
4 of information among the respective switching nodes 11(n). The communication links 13(p) may
5 utilize any convenient information transmission medium, including, for example, wires for carrying
6 electrical signals, optical fiber links for carrying optical signals, and so forth. Each communication
7 link 13(p) is preferably bi-directional, allowing the switching nodes 11(n) to transmit and receive
8 signals among each other and with customer premises equipment 12(m) connected thereto over the
9 same link; depending on the particular type of medium selected for respective communication links
10 13(p), multiple media may be provided for transferring signals in opposite directions thereby to
11 provide the bidirectional link. In addition, each communication link 13(p) depicted in FIG. 1 may
12 comprise a plurality of physical uni- or bi-directional links to provide redundancy in the event of a
13 failure in one of the links or transmitting or receiving circuitry associated therewith.

14 Data is transferred in the network 10 in the form of packets. Generally, a packet includes a
15 header portion and a data portion. The header portion includes information that assists in routing
16 the packet through the network, with the specific information depending on the particular packet
17 routing protocol that is used in routing packets through the network. In connection with network 10,
18 any of a number of well-known packet routing protocols may be used. In one embodiment, which
19 will be described in connection with FIG. 2, a connectionless routing protocol is used, in particular,
20 the well-known Internet protocol ("IP"). A second embodiment, in which a connection-oriented
21 routing protocol, in which packets are transferred over virtual circuits, will further be described.
22 In the connectionless routing protocol, the header portion typically includes address information
23 including a source address identifying the particular source device 12(mS) that generated the packet
24 and a destination address identifying the particular destination device 12(mD) that is to receive the
25 packet. In the IP protocol, a packet may be of variable length and the header typically will also
26 include length information to identify the length of the packet. The header also typically includes
27 other information, including, for example, protocol identifier information that identifies the

1 particular protocol that defines the structure of the packet. The data portion contains the data
2 payload of the packet. The packet may also include, as part of the header portion, data portion or
3 otherwise, error detection information that can be used to determine whether an error occurred in
4 transferring the packet.

5 A source device 12(mS), after generating a packet for transfer to a destination device
6 12(mD), will provide the packet to the switching node 11(n) connected thereto. The switching node
7 11(n) will use the destination address in the packet to attempt to identify a "route," which associates
8 a destination address with one of the communication links 13(p) connected thereto over which it is
9 to transfer the packet to forward it (that is, the packet) to either the destination device 12(mD), if the
10 switching node 11(n) is connected to the destination device 12(mD), or to another switching node
11 11(n') (n'≠n) along a path to the destination device 12(mD). Each switching node maintains a
12 network topology database, which identifies the switching nodes 11(n) comprising the network, the
13 communication links 13(p) which interconnect them, and other information such as communications
14 "cost" information associated with each communication link 13(p), which they (that is, the switching
15 nodes 11(n)) use to determine the route information. Several methodologies are known for
16 generating route information from network topology information in the network topology database,
17 including the well-known OSPF ("open shortest path first") methodology. If the switching node can
18 identify a route for the received packet, it will forward the packet over the communication link
19 identified by the route. Each switching node 11(n'), 11(n''),..., which receives the packet will
20 perform a similar operation. If all of the switching nodes have respective routes for the destination
21 address, the packet will eventually arrive at the destination device 12(mD).

22 The invention provides an arrangement, which will be described in connection with the
23 network 10 depicted in FIG. 1, for facilitating recovering from failures or other malfunctions in
24 connection with a communication link 13(p) interconnecting two switching nodes 11(n). A
25 switching node 11(n) and a device 12(m) may determine that a failure or other malfunction has
26 occurred using any of a number of conventional arrangements. For example, elements connected
27 to a particular communication link 13(p) can verify that it has not failed or otherwise malfunctioned

1 by, for example, periodically transferring link status packets thereover, particularly if other packets
2 are not being transferred thereover. In that case, if a switching node 11(n) or device 12(m) does not
3 receive a link status packet over a communication link 13(p) within a predetermined window of time
4 following receipt of a previous link status packet, and if other types of packets have not been
5 received in the meantime, the respective switching node 11(n) or device 12(m) can determine that
6 the communication link 13(p) has failed or otherwise malfunctioned. If a communication link 13(p)
7 is a uni-directional communication link, in which packets are transferred from one switching node
8 11(n) to another switching node 11(n') ($n' \neq n$), but not from switching node 11(n') to switching node
9 11(n), link status packets can still be used to verify that the communication link 13(p) has not failed
10 or otherwise malfunctioned. In that case, the switching node 11(n) can transfer link status packets
11 over the communication link 13(p), and the switching node 11(n') can either notify the switching
12 node 11(n) that it (that is, switching node 11(n')) has, or has not, received the link status packet over
13 another path therebetween through the network, which path may ~~extent~~ through another switching
14 node 11(n'') ($n'' \neq n', n$). Other mechanisms by which the switching nodes and devices can determine
15 whether a communication link 13(p) is operating properly, or whether it has failed or otherwise
16 malfunctioned, including, for example, mechanisms include detecting the loss of physical signal on
17 the communication link or mechanisms similar to those which are implemented in a network
18 constructed using the well-known SONET, SDH and TDM protocols, will be apparent to those
19 skilled in the art.

20 As noted above, The invention provides an arrangement for facilitating recovering from
21 failures or other malfunctions in connection with a communication link 13(p) interconnecting two
22 switching nodes 11(n). In the first embodiment, in which the network 10 transfers packets in a
23 connectionless manner, each switching node 11(n) establishes, for each communication link 13(p)
24 which interconnects it (that is, the switching node 11(n)) and another switching node 11(n'), a
25 permanent, pre-established special-purpose "bypass" virtual circuit which bypasses the
26 communication link, over which the switching node 11(n) can transfer packets in the event of a
27 failure or other malfunction in connection with the communication link 13(p). Thus, if a

1 communication link 13(p) fails or otherwise malfunctions, and if the switching node 11(n) receives
2 a packet which its routing table indicates would be properly transferred over that communication link
3 13(p), instead of discarding the packet or routing it in a connectionless manner to another switching
4 node 11(n") (which may result in the other switching node 11(n') returning the packet to the
5 switching node 11(n)), the switching node 11(n) embeds the packet in one or more virtual circuit
6 packets and transfers it over the pre-established bypass virtual circuit to the switching node 11(n').
7 When the switching node 11(n') receives the virtual circuit packet(s) from the switching node 11(n)
8 over the bypass virtual circuit, it will extract the original packet from the virtual circuit packet(s) and
9 forward it toward the destination device 12(mD) in the usual manner.

10 As a specific illustration, suppose the switching node 11(1) receives a packet (from, for
11 example, the device 12(1)), for transfer to the device 12(2). Suppose further that the routing table
12 of switching node 11(1) indicates that packets to be transferred to device 12(2) are to be transferred
13 thereby (that is, by switching node 11(1)) over communication link 13(10) to switching node 11(2)),
14 but that that communication link 13(10) has failed or otherwise malfunctioned. Finally, suppose that
15 the switching nodes 11(1) and 11(2) have established a bypass virtual circuit through switching node
16 11(3), including communication links 13(9) and 13(11). In that case, the switching node 11(1), after
17 receiving the packet for transfer to the device 12(2), will embed the packet in one or more virtual
18 circuit packets for transfer over the bypass virtual circuit, and will transfer the virtual circuit packets
19 over the communication link 13(9) to switching node 11(3). The switching node 11(3), as it receives
20 the virtual circuit packets, will forward them to the switching node 11(2) over communication link
21 13(11). The switching node 11(2) will reconstruct the original packet from the virtual circuit
22 packet(s) either as it receives the virtual circuit packets or after it has received all of them. In any
23 case, after the switching node 11(2) has constructed the original packet, it will forward the packet
24 to its intended destination, in this case device 12(2).

25 Similar operations can occur in connection with, for example, packets received by switching
26 node 11(1) for transfer to device 12(3) or 12(4), connected to switching node 11(4), if its (that is,
27 switching node 11(1)'s) routing table indicates that such packets are to be transferred over

1 communication link 13(10) and through switching node 11(2), and if communication link 13(10) has
2 failed or otherwise malfunctioned. Thus, if the switching node 11(1) receives a packet for transfer
3 to device 12(3), and if communication link 13(10) has failed or otherwise malfunctioned, the
4 switching node 11(1), after receiving the packet, will embed the packet in one or more virtual circuit
5 packets for transfer over the bypass virtual circuit, and will transfer the virtual circuit packets over
6 the communication link 13(9) to switching node 11(3). The switching node 11(3), as it receives the
7 virtual circuit packets, will forward them to the switching node 11(2) over communication link
8 13(11). The switching node 11(2) will reconstruct the original packet from the virtual circuit
9 packet(s) either as it receives the virtual circuit packets or after it has received all of them. In any
10 case, after the switching node 11(2) has constructed the original packet, it will forward the packet
11 to switching node 11(4). The switching node 11(4), after it receives the original packet, can forward
12 it to the device 12(4) as destination.

13 Bypass virtual circuits can be established in network 10 for each communication link 13(p)
14 in a similar manner. If a communication link 13(p) is a uni-directional communication link, the
15 respective bypass virtual circuit therefor can be a uni-directional bypass virtual circuit. On the other
16 hand, if a communication link 13(p) is a bi-directional communication link, the respective bypass
17 virtual circuit therefor can be a bi-directional bypass virtual circuit, or alternatively the switching
18 nodes 11(n), 11(n') can establish two uni-directional bypass virtual circuits therefor. The virtual
19 circuits can be established to transfer virtual circuit packets in accordance with any connection-
20 oriented packet transfer methodology, such as MPLS, ATM, or the like.

21 As noted above, each switching node 11(n) includes a network topology database, from
22 which it determines route information to be used in routing packets over the communication links
23 13(p) connected thereto. The switching nodes 11(n) can establish and update information in the
24 network topology database in a conventional manner by broadcasting link state advertising messages
25 through the network 10. Each link state advertising message identifies the switching node 11(n),
26 the communication links 13(p) connected thereto, and the communication cost information that is
27 applicable to each of the communication links 13(p). After a switching node 11(n) receives a link

1 state advertising message, it will determine whether the link state advertising message indicates that
2 the network topology has changed sufficiently to warrant updating the route information which it
3 maintains in its routing tables, and, if so, will update the route information. If a communication link
4 13(p) fails or otherwise malfunctions, if the switching nodes 11(n), 11(n') connected thereto
5 broadcast link state advertising messages over the network so indicating, since the link state
6 advertising messages will be received by the other switching nodes and processed at different times,
7 their network topology databases and routing tables will generally be inconsistent with each other
8 for at least some time following the determination that the communication link 13(p) has failed or
9 otherwise malfunctioned. It will be appreciated that, by use of the bypass virtual circuits, the
10 switching nodes 11(n), 11(n') may avoid the necessity of notifying the other switching nodes that
11 the communication link has failed or otherwise malfunctioned. Alternatively, the switching nodes
12 11(n), 11(n') may provide a notification that the communication cost related to transfers over the
13 virtual circuit has increased, particularly if the total communication cost related to the
14 communication links forming the path of the bypass virtual circuit is higher than that for the
15 communication link which has failed or otherwise malfunctioned and is bypassed by the respective
16 bypass virtual circuit.

17 The embodiment described above allows the network to operate correctly even in the event
18 of a failure of a communication link 13(p), and therefore eliminates the need to immediately notify
19 all of the switching nodes 11(n) in the event of a link failure. Such a notification ensures that the
20 network will relatively rapidly converge on correct routes which take the link failure into account.
21 However, in a connectionless network using conventional prior art routing methods, such
22 notification also ensures that there will be a short period of time during which the switching nodes
23 11(n) have an inconsistent view of the topology of the network, potentially resulting in looping
24 packets, lost packets, or other network disruption. However, in some cases it might be desirable to
25 provide a notification to the other switching nodes that a communication link has failed or otherwise
26 malfunctioned so that the other switching nodes can update their network topology databases and
27 route information accordingly. In that case, the switching nodes 11(n), 11(n') can, for example,

(i) notify a network manager, who can determine the cause of the failure or malfunction, and whether recovery will require enough time to warrant updating the network topology databases and route information generally throughout the network. If the network manager determines that the recovery can be performed quickly enough that updates are not warranted, he or she may merely proceed with the recovery, without enabling updates to the network topology databases or route information. On the other hand, if the network manager determines that the recovery can not be performed quickly enough that updates are not warranted, he or she may, in addition to proceeding with the recovery, enable updates to the network topology databases of the respective switching nodes; in that case, if the switching nodes can determine whether the update warrants updating the route information in their respective routing tables and, if so, perform the updates.

(ii) generate link state advertising messages for broadcast through the network, but specify that the switching nodes only update the route information in their routing tables at a predetermined point in time. The predetermined point in time may be selected so as to guarantee that the network 10 is not busy. Since all of the switching nodes will apply the update at the predetermined point in time, their routing tables will not be inconsistent.

(iii) initially broadcast link state advertising messages which do not indicate that the communication link has failed or otherwise malfunctioned, but instead that the communication cost associated therewith has increased by a selected small amount. After an amount of time, which is selected to be at least the amount of time required for the last switching node to receive the link state advertising message, the switching nodes 11(n), 11(n') will again broadcast link state advertising messages for which the communication cost associated with the communication link is similarly increased. As the other switching nodes receive the link state advertising messages, they will update their routing tables as warranted by the increased communication cost for the respective communication link. These operations will be repeated through a plurality of iterations until, eventually, the communication cost has reached a level such that no network traffic is computed (using the OSPF methodology, as described above) to use that communication link. At that point, the switching nodes 11(n), 11(n') can broadcast link state advertising messages indicating that the

1 communication link has failed.

2 It will be appreciated that, after a communication link 13(p) has failed or otherwise malfunctioned,
3 and the bypass virtual circuit used in substitution therefor, the communication cost associated with
4 the communication link 13(p), and with the communication links 13(p'), 13(p'') (p', p'' ≠ p) which are
5 used for the bypass virtual circuit, can be advertised in the link state advertising messages as having
6 increased, since there will be increased congestion associated therewith. The other switching nodes
7 can use the increased communication cost as warranted in updating their route information, which
8 can help balance the traffic across communication links in the network. After the communication
9 link 13(p) has been repaired, the link state advertising messages can indicate a corresponding
10 reduction in communication cost associated therewith and with the communication links 13(p'),
11 13(p'') used for the bypass virtual circuit, and the other switching nodes can use the decreased
12 communication cost as warranted in updating their route information, which can restore the network
13 to its previous traffic condition.

14 If there are widespread failures or more than one simultaneous failure, causing a link to fail
15 at the same time that the associated bypass virtual circuit or circuits, it will generally be desirable
16 to provide notifications to the various switching nodes 11(n) using conventional routing protocol
17 mechanisms.

18 With this background, operations performed by a switching node 11(n) in connection with
19 this aspect of the invention will be described in connection with FIG. 2. With reference to FIG. 2,
20 the switching node 11(n) performs a plurality of operations to establish bypass virtual circuits for
21 each of the communication links 13(p) interconnecting it (that is, the switching node 11(n)) and other
22 switching nodes 11(n'), 11(n''),..., in the network 10. In that operation, the switching node 11(n) will
23 determine whether it is connected to a communication link 13(p) for which no bypass virtual circuit
24 exists (step 100). If the switching node 11(n) makes a positive determination in step 100, it will
25 identify the other switching node 11(n'), 11(n''),... that is connected to that communication link 13(p)
26 (step 101), and identify a path through the network to that other switching node (step 102). The

1 switching node 11(n) can use any conventional methodology for identifying a path through the
2 network to the other switching node in step 102, including the OSPF methodology. After the
3 switching node 11(n) has identified the path, it will enable the bypass virtual circuit to be established
4 over the identified path (step 103). In that operation, any conventional virtual circuit establishment
5 methodology can be used to establish the bypass virtual circuit. However, in connection therewith,
6 at least the switching node 11(n) and the other switching node to which the bypass virtual circuit is
7 being established will be aware that the virtual circuit is a bypass virtual circuit, so that they will,
8 when a packet is received for transfer over the communication link 13(p), if the communication link
9 has failed or otherwise malfunctioned, instead of transferring the packet over the communication link
10 13(p), embed the packet in one or more virtual circuit packets, transfer the virtual circuit packets
11 over the bypass virtual circuit, and extract the original packet therefrom for transfer to the
12 destination. In any case, following step 103, the switching node 11(n) will return to step 100, and
13 repeat steps 100 through 103 through one or more iterations until a bypass virtual circuit has been
14 established for each communication link connected thereto, after which it can exit (step 104).

15 After the bypass virtual circuit has been established for a communication link 13(p), the
16 switching node 11(n) and other switching node connected thereto can use it in the case of a failure
17 or other malfunction of the communication link 13(p). In that connection, when the switching node
18 11(n) receives a packet for transfer over the communication link 13(p) (step 110), it will initially
19 determine whether the communication link 13(p) is operational (step 111). If the switching node
20 11(n) makes a positive determination in step 111, it will transfer the packet over the communication
21 link 13(p) (step 112). On the other hand, if the switching node 11(n) makes a negative determination
22 in step 111, which will occur if the communication link 13(p) has failed or otherwise malfunctioned,
23 it will embed the packet in one or more virtual circuit packets (step 113) and transfer it or them over
24 the bypass virtual circuit (step 114). When the other switching node receives the virtual circuit
25 packets, it will extract the original packet therefrom (step 115). Following step 112 or step 115, the
26 other switching node will transfer the packet (that is, the original packet) toward its intended
27 destination (step 116).

1 As noted above, in connection with a second embodiment, in which the network 10 transfers
2 packets in a connection-oriented manner, packets are normally transferred through the network using
3 constructs such as virtual circuits. In that embodiment, each packet, instead of providing a
4 destination identifier identifying the actual destination for the packet, provides a virtual circuit
5 identifier which the switching node that receives the packet uses to identify the communication link
6 over which the packet is to be transferred, using information from a virtual circuit table. In any case,
7 in the second embodiment, switching nodes 11(n) comprising the network 10 also establish bypass
8 virtual circuits which they use if a communication link 13(p) fails or otherwise malfunctions, in
9 connection with transfer of packets for virtual circuits for which the communication link 13(p) forms
10 part of the path. However, in the case of the second embodiment, the switching nodes 11(n)
11 establish either

12 (i) establish a bypass virtual circuit for each virtual circuit for which the communication link
13 forms part of the path, which will be used in connection with packets associated with the respective
14 virtual circuit if the communication link 13(p) fails or otherwise malfunctions, or

15 (ii) establish one or several virtual circuits over which virtual circuit packets ~~associate~~ ^{associated} with
16 some subset of the virtual circuit(s) for which the communication link forms part of the path, which
17 will be used in connection with packets associated with virtual circuit(s) in the respective subset, if
18 the communication link fails or otherwise malfunctions.

19 It will be appreciated that item (ii) above will find utility primarily in connection with virtual circuit
20 transfer methodologies which allow packets from multiple virtual circuits to be merged into a single
21 virtual circuit for transfer. In connection with either item (i) or item (ii), it will be appreciated that
22 the bypass virtual circuit established by one switching node 11(n) need not terminate at the switching
23 node 11(n') connected to the communication link, but instead may terminate at another switching
24 node 11(n'') downstream thereof along the path for the virtual circuit, if, for example, the switching
25 node 11(n) determines that that will enhance efficient transfer of packets along the virtual circuit.

26 In either case, when a switching node 11(n) receives a packet associated with a virtual

1 circuit for which a communication link 13(p) which has failed or otherwise malfunctioned forms part
2 of the path, it (that is, the switching node 11(n)) will transfer the packet over the bypass virtual
3 circuit, in a manner similar to that described above in connection with the connectionless transfer
4 methodology. Similarly, when the switching node which forms the destination end of the bypass
5 virtual circuit receives a packet therefrom, it will resume transfer of the packet over the virtual circuit
6 (that is, the virtual circuit whose path included the communication link 13(p) which had failed or
7 otherwise malfunctioned.)

8 The invention provides a number of advantages. In particular, the invention provides for the
9 rapid recovery in connection with transfer of packets through a network in the event of a failure or
10 other malfunction in connection with a communication links, while minimizing problems that can
11 occur if the network topology information and routing information maintained by the respective
12 switching nodes comprising the network is inconsistent.

13 It will be appreciated that numerous modifications may be made to the invention. For
14 example, although the invention has been described as making use of virtual circuits to provide the
15 mechanism for bypassing a communication link which has failed or otherwise malfunctioned, it will
16 be appreciated that other arrangements, such as source routing, in which the switching node 11(n)
17 connected to a communication link which has failed or otherwise malfunctioned determines the route
18 to the other switching node 11(n') through one or more other switching nodes. In that case, the
19 switching node 11(n) would embed the packet in a source-routed packet for transfer to the other
20 switching node 11(n'). The switching node 11(n'), when it receives the source-routed packet, extracts
21 the original packet for transfer towards the destination.

22 In addition, it will be appreciated that any convenient virtual circuit transfer methodology
23 can be used, including ATM, MPLS, etc., for use in connection with bypass virtual circuits.

24 In addition, although the invention has been described in connection with providing a bypass
25 virtual circuit for use in the event of failure of a communication link 13(p), it will be appreciated that
26 bypass virtual circuits can similarly be established to bypass other components, including, for

1 example, a switching node, in the event of a failure or other malfunction of the switching node. For
2 example, with reference to FIG. 1, if switching node 11(2) wishes to establish a bypass virtual circuit
3 to be used in the event of a failure of switching node 11(3) for transfer of packets to switching node
4 11(6), it can establish a bypass virtual circuit to node 11(6) through switching node 11(1). Thus, if
5 switching node 11(2) determines that switching node 11(3) has failed or otherwise malfunctioned,
6 it will transfer packets to switching node 11(6) over that bypass virtual circuit as described above.
7 Bypass virtual circuits for use in the event a switching node 11(n) fails or otherwise malfunctions
8 may be established either instead of or in addition to the bypass virtual circuits established for use
9 in the event of a failure or other malfunction of a communication link 13(p).

10 It should be noted that, in general it may be difficult to rapidly determine whether a
11 communication link 13(p) or switching node 11(n) has failed or otherwise malfunctioned,
12 particularly if such determination requires transfer of packets through other switching nodes. For
13 example, with reference to FIG. 1, if switching node 11(2) detects a failure in connection with
14 communication link 13(11), the switching node 11(2) will not be immediately able to determine
15 whether the communication link 13(p) itself has failed or whether the switching node 11(3)
16 connected thereto has failed. However, bypass virtual circuits can be set up to other nodes in the
17 network which bypass multiple network components, such that the failure can be bypassed
18 regardless of which if the network components has failed.

19 Furthermore, it will be appreciated that bypass virtual circuits can be established in a network
20 for all, or any subset of communication links and/or switching nodes in the network.

21 In addition, it will be appreciated that the invention can be used in connection with networks
22 carrying any type of digital data, including voice or video data, as well as data for use in digital
23 processing.

24 In addition, although the switching nodes 11(n) have been described as being similar to the
25 switching node described in the aforementioned Schwartz, et al., application, it will be appreciated
26 that the invention may be used in connection with switching nodes of any structure or architecture.

1 It will be appreciated that a system in accordance with the invention can be constructed in
2 whole or in part from special purpose hardware or a general purpose computer system, or any
3 combination thereof, any portion of which may be controlled by a suitable program. Any program
4 may in whole or in part comprise part of or be stored on the system in a conventional manner, or it
5 may in whole or in part be provided in to the system over a network or other mechanism for
6 transferring information in a conventional manner. In addition, it will be appreciated that the system
7 may be operated and/or otherwise controlled by means of information provided by an operator using
8 operator input elements (not shown) which may be connected directly to the system or which may
9 transfer the information to the system over a network or other mechanism for transferring
10 information in a conventional manner.

11 The foregoing description has been limited to specific embodiments of this invention. It will
12 be apparent, however, that various variations and modifications may be made to the invention, with
13 the attainment of some or all of the advantages of the invention. It is the object of the appended
14 claims to cover these and such other variations and modifications as come within the true spirit and
15 scope of the invention.

16 What is claimed as new and desired to be secured by Letters Patent of the United States is: